arXiv:1704.02065v1 [cs.DM] 7 Apr 2017

Computational Approaches for Zero Forcing and Related Problems

Boris Brimkov, Caleb C. Fast, and Illya V. Hicks

Department of Computational & Applied Mathematics, Rice University, Houston, TX 77005, USA
boris.brimkov@rice.edu, caleb.c.fast@rice.edu, ivhicks@rice.edu

Abstract. In this paper, we propose computational approaches for the zero forcing problem, the connected zero forcing problem, and the problem of forcing a graph within a specified number of timesteps. Our approaches are based on a combination of integer programming models and combinatorial algorithms, and include formulations for zero forcing as a dynamic process, and as a set-covering problem. We explore several solution strategies for these models, and numerically compare them to the well-known Wavefront algorithm for zero forcing developed by Grout et al. Our proposed algorithms for connected zero forcing and for controlling the propagation time are the first general-purpose computational methods for these problems; they are comparable in performance to Wavefront and superior to brute force computation.

Keywords: Zero forcing, connected zero forcing, propagation time, integer programming, set-covering

1 Introduction

Zero forcing is an iterative graph coloring process where at each discrete time step, a colored vertex with a single uncolored neighbor forces that neighbor to become colored. A zero forcing set of a graph is a set of initially colored vertices which forces the entire graph to become colored. The zero forcing number is the cardinality of the smallest zero forcing set. Zero forcing was initially introduced to bound the maximum nullity of the family of symmetric matrices described by a graph [2]; it was also independently studied in quantum physics [15] and theoretical computer science [49], and has since found a variety of uses in physics, logic circuits, coding theory, power network monitoring, and in modeling the spread of diseases and information in social networks; see [7,15,17,35,36,44,50] for more details.

Connected zero forcing is a variant of zero forcing in which the initially colored set of vertices induces a connected subgraph. The connected zero forcing number of a graph is the cardinality of the smallest connected set of initially colored vertices which forces the entire graph to be colored (i.e., the smallest connected zero forcing set). Applications and various structural and computational aspects of connected zero forcing have been investigated in [11,12,13]. Other

variants of zero forcing, such as positive semidefinite zero forcing [8,26,37,46], fractional zero forcing, signed zero forcing [34], and k-forcing [3,39] have also been studied. These are typically obtained by modifying the zero forcing color change rule, or adding certain restrictions. The number of timesteps in the zero forcing process after which a graph becomes colored is also a problem of interest (see, e.g., [10,16,22,36,46]). Connected variants of other graph problems – such as connected domination and connected power domination [20,24,29,33] – have been extensively studied as well.

Computing the zero forcing number and connected zero forcing number of a graph are both NP-complete problems [1,11,49]; nevertheless, it is important to develop practical algorithms for solving these problems, at least on moderately-sized graphs. The state-of-the-art approach for computing the zero forcing number of a graph is a combinatorial algorithm called Wavefront, developed by Grout et al. [19]. While this algorithm works well for the zero forcing problem, it is not flexible and cannot accommodate additional constraints, such as assuring connectivity of the solution or limiting the propagation time. A lot of effort has been put into developing closed formulas, efficient algorithms, characterizations, and bounds for the zero forcing numbers of graphs with special structure (see, e.g., [2,9,11,12,25,27,38,40]), but relatively little progress has been made on developing computational methods for general graphs.

1.1 Main Contributions

In this paper, we explore approaches for computing the zero forcing number of a graph using integer programming. In particular, we present formulations of zero forcing based on two different perspectives – one as a dynamic process, and the other as a set-covering problem. We explore several solution strategies of these models – such as direct computation, constraint generation, and generation of facet-inducing constraints – and compare their performance to Wavefront on different types of random graphs.

We also propose a combinatorial algorithm for computing the connected zero forcing number of a graph, and extend the proposed integer programming models to the connected zero forcing problem by adding connectivity constraints. In doing so, we explore several different types of connectivity constraints which have been used in problems like connected domination, Steiner trees, and forest planning. Until now, there have not been any computational approaches for connected zero forcing of general graphs other than brute-force computation.

Finally, we adapt one of our integer programs to find zero forcing sets which force the graph within a specified number of timesteps, and have minimum cardinality among all such sets. To our knowledge, there have not been any previously-implemented algorithms for this problem.

Our computational experiments show that the Wavefront algorithm generally outperforms the integer programming models for zero forcing. However, for connected zero forcing, our proposed integer programming models were mutually competitive, and significantly outperformed the combinatorial brute force and branch and bound algorithms. Moreover, these approaches were faster, and

able to handle larger graphs, than the Wavefront algorithm and the zero forcing analogues of the integer programs. This is somewhat surprising, since the connected variants of problems like domination and power domination have typically proven more difficult to solve computationally, due to their non-locality (see [33] for more details). Since the connected zero forcing number is an upper bound to the zero forcing number, the proposed approaches for connected zero forcing can be used to obtain upper bounds or approximations to the zero forcing number, especially for graphs which are too large for Wavefront.

The paper is organized as follows. In the next section, we recall some graph theoretic notions, specifically those related to zero forcing. In Section 3, we present combinatorial approaches for computing the zero forcing number and connected zero forcing number of a graph; in Section 4, we present integer programming approaches for these problems. In Section 5, we describe the implementation of our proposed approaches, and compare them through computational experiments on various types of random graphs. We conclude with some final remarks and open questions in Section 6.

2 Preliminaries

A graph $G = (V, E)$ consists of a vertex set V and an edge set E of two-element subsets of V. In this paper, we consider *simple* graphs, for which a subset $\{v, w\} \in E$ must have $v \neq w$ and E contains at most one copy of $\{v, w\}$. The *order* and *size* of G are denoted by $n = |V|$ and $m = |E|$, respectively. Two vertices $v, w \in V$ are *adjacent*, or *neighbors*, if $\{v, w\} \in E$. The *neighborhood* of $v \in V$ is the set of all vertices which are adjacent to v, denoted $N(v)$; the *closed neighborhood* of v, denoted $N[v]$, is the set $N(v) \cup \{v\}$. Similarly, given $S \subset V$, $N(S)$ denotes the union of the neighborhoods of the vertices in S, and $N[S]$ denotes the set $N(S) \cup S$. The *degree* of $v \in V$ is defined as $d(v) = |N(v)|$. Given $S \subset V$, the *induced subgraph* $G[S]$ is the subgraph of G whose vertex set is S and whose edge set consists of all edges of G which have both endpoints in S. For other graph theoretic terminology and definitions, we refer the reader to [48].

Given a graph $G = (V, E)$ and a set $S \subset V$ of initially colored vertices, the *color change rule* dictates that at each integer-valued time step, a colored vertex u with a single uncolored neighbor v *forces* that neighbor to become colored. The *closure* of S, denoted $cl(S)$, is the set of colored vertices obtained after the color change rule is applied until no new vertex can be forced; it can be shown that the closure of S is uniquely determined by S (see [2]). A *zero forcing set* is a set whose closure is all of V; the *zero forcing number* of G, denoted $Z(G)$, is the minimum cardinality of a zero forcing set. A *chronological list of forces* associated with a zero forcing set Z is a sequence of forces applied to obtain the closure of Z in the order they are applied. A *forcing chain* for a chronological list of forces is a maximal sequence of vertices $(v_1, \ldots, v_k)$ such that v_i forces v_{i+1} for $1 \leq i \leq k-1$. Each forcing chain induces a distinct path in G, one of whose endpoints is an initially colored vertex; the other is called a *terminal*

vertex. A *fort*, defined by Fast and Hicks [30], is a non-empty set $F \subset V$ such that no vertex outside F is adjacent to exactly one vertex in F. A zero forcing set *restrained* by $S \subset V$ is a zero forcing set which contains S; $Z(G;S)$ denotes the cardinality of the smallest zero forcing set restrained by S (cf. [13]).

A *connected zero forcing set* of G is a zero forcing set of G which induces a connected subgraph. The *connected zero forcing number* of G, denoted $Z_c(G)$, is the cardinality of a minimum connected zero forcing set of G. For short, we may refer to these as *connected forcing set* and *connected forcing number*. Note that a disconnected graph cannot have a connected forcing set.

An *empty* graph is a graph with no edges. A graph is *cubic* if all of its vertices have degree 3. Let $C(n,k)$ be the graph with vertex set $\{0,\ldots,n-1\}$ and edge set $\{\{i,j\} : 0 \leq i,j \leq n-1, |i-j| \leq k/2\}$. A *Watts-Strogatz* graph with parameters n, k, and β refers to a graph obtained from $C(n,k)$ by replacing each edge of $C(n,k)$ with probability β by a randomly chosen edge. When n is clear from the context (or for a predetermined set of values of n), we will refer to the Watts-Strogatz graphs with parameters n, k, and β as $\mathrm{WS}(k,\beta)$. Watts-Strogatz graphs were introduced in [47], and are a popular random graph model; they are meant to have small-world properties such as short average path lengths and high clustering.

Finally, we will use the notation $[n]$ to represent the set $\{1,\ldots,n\}$.

3 Combinatorial Approaches

In this section, we describe several combinatorial approaches for computing the zero forcing and connected forcing numbers of a graph $G = (V,E)$. The trivial approach for finding a minimum zero forcing set of G is to iteratively compute the closures of all subsets of V of size i, starting from $i = 1$ and incrementing i, until a zero forcing set is found. Similarly, to find a minimum connected forcing set, one could again generate subsets of vertices of increasing size, check whether each set induces a connected subgraph, and stop when the first connected set whose closure is V is found.

This brute force approach works well when the graph is known *a priori* to have a very small or very large forcing number (in the latter case, one would start from $i = n$, decrement i as soon as a forcing set is found, and stop when all sets of vertices of a certain size are not forcing). Similarly, the brute force approach can be used in conjunction with theoretical bounds on the forcing number in terms of other efficiently-computable parameters (see, e.g., [2,9,11,12,43]). In particular, if it is determined that $k_1 \leq Z_c(G) \leq k_2 < \frac{n}{2}$, it can be checked whether each of the $\binom{n}{k_1} + \cdots + \binom{n}{k_2}$ sets of vertices of appropriate size is connected and forcing in $O(n^2)$ time, so $Z_c(G)$ can be computed in $O((k_2 - k_1)n^{2+k_2})$ time (the same applies to $Z(G)$). Other advantages of the brute force approach are that it is easy to implement, uses little memory, and can be easily parallelized, since closures of different sets of vertices can be computed independently.

Nevertheless, in practice, the brute force algorithm is usually outperformed by the other algorithms discussed in the sequel. Section 3.1 describes the Wave-

front algorithm – a dynamic programming style improvement of the brute force algorithm, which stores minimum forcing sets of certain subgraphs of G and uses them to build minimum forcing sets of larger subgraphs. Thus, it avoids checking all possible subsets of vertices at the expense of increased memory. Section 3.2 gives a branch-and-bound style improvement of the brute force algorithm for connected forcing; instead of generating all subsets of vertices and checking whether they are connected and forcing, this algorithm generates only connected subgraphs, checks whether they are forcing, and prunes the search tree based on the best zero forcing set found.

3.1 Wavefront Algorithm

In this section, we give a description of the combinatorial algorithm for zero forcing known as Wavefront, developed by Grout et al. [18]. To our knowledge, this algorithm is the only previously-implemented computational method for the zero forcing problem (aside from brute force), and a proof of its correctness does not appear elsewhere in print. We prove that the Wavefront algorithm is correct in Theorem 1 and give a result about its worst-case memory requirements in Theorem 2.

Algorithm 1: Wavefront Algorithm [18]

```
Data: G = (V, E)
Result: Zero forcing number of G
𝒞 = {(∅, 0)};
for R ∈ [n] do
    for (S, r) ∈ 𝒞 do
        for v ∈ V do
            k = number of uncolored neighbors of v;
            C = cl(S ∪ N[v]);
            if (C, i) ∉ 𝒞 for i ≤ R then
                if v ∉ S and k ≤ R − r then add (C, r + k) to 𝒞;
                if v ∈ S and k − 1 ≤ R − r then add (C, r + k − 1) to 𝒞;
    if (V, z) ∈ 𝒞 then return z;
```

Lemma 1. *Let F be a fort of $G = (V, E)$ and $S \subset V$ be a set that does not contain any set H consisting of a vertex $v \in N[F]$ together with all-but-one vertices in $N(v) \cap F$. Then $F \backslash S$ is a fort of G.*

Proof. Since F is a fort, every vertex in $N[F]$ is either in F or has at least two neighbors in F. Let u be a vertex in $N[F \backslash S]$, and hence also in $N[F]$. Since S does not contain any vertex in $N[F]$ together with all-but-one of its neighbors in F, u must either not be in S or u must have at least two neighbors that are

in F but not in S. Thus, u must either be in $F\backslash S$ or u must have at least two neighbors in $F\backslash S$. Thus, $F\backslash S$ is a fort of G. □

Corollary 1. *Let F be a fort of $G = (V, E)$. Then, any zero forcing set of G must contain some set H consisting of a vertex $v \in N[F]$ together with all-but-one vertices in $N(v) \cap F$.*

Proof. Suppose for contradiction that some zero forcing set S of G does not contain any set H consising of a vertex $v \in N[F]$ together with all-but-one vertices in $N(v) \cap F$. By Lemma 1, $F\backslash S$ is a fort of G. Thus, S cannot force any vertex in $F\backslash S$, contradicting that S is a zero forcing set of G. □

Theorem 1. *The Wavefront algorithm returns $Z(G)$.*

Proof. Observe that for any non-forcing set S, $V\backslash cl(S)$ is a fort of $G = (V, E)$. Let $H_0 = \emptyset$, and note that the algorithm is initialized with $(cl(H_0), |H_0|)$ being added to $\mathcal{C}$.

Let Z_0 be a set that realizes $Z(G; H_0)$; since $H_0 = \emptyset$, it follows that $|Z_0| = Z(G)$. Moreover, since $H_0 = \emptyset$, $V\backslash cl(H_0)$ is a fort. Thus, by Corollary 1, Z_0 must contain a set H_1 consisting of some vertex $v_1 \in N[V\backslash cl(H_0)]$ together with all-but-one of its neighbors which are not in $cl(H_0)$. Since Z_0 contains H_0 and H_1, there cannot be a set S with $|S| < |H_0 \cup H_1|$ and $cl(H_0 \cup H_1) \subset cl(S)$ (otherwise, $(Z_0\backslash(H_0 \cup H_1)) \cup S$ would be a zero forcing set of size smaller than $Z(G)$). Therefore, $(cl(H_0 \cup H_1), |H_0 \cup H_1|)$ must be added to $\mathcal{C}$ by step $|H_0 \cup H_1|$ of the algorithm.

Let Z_1 be a set that realizes $Z(G; H_0 \cup H_1)$; since $H_0 \cup H_1$ is contained in some minimum zero forcing set (namely Z_0), it follows that $|Z_1| = Z(G)$. If $cl(H_0 \cup H_1) \neq V$, then $V\backslash cl(H_0 \cup H_1)$ is a fort. Thus, by Corollary 1, Z_1 must contain a set H_2 consisting of some vertex $v_2 \in N[V\backslash cl(H_0 \cup H_1)]$ together with all-but-one of its neighbors which are not in $cl(H_0 \cup H_1)$. Since Z_1 contains $H_0 \cup H_1$ and H_2, there cannot be a set S with $|S| < |H_0 \cup H_1 \cup H_2|$ and $cl(H_0 \cup H_1 \cup H_2) \subset cl(S)$ (otherwise, $(Z_1\backslash(H_0 \cup H_1 \cup H_2)) \cup S$ would be a zero forcing set of size smaller than $Z(G)$). Therefore, since $(cl(H_0 \cup H_1), |H_0 \cup H_1|)$ is in $\mathcal{C}$ by step $|H_0 \cup H_1|$, $(cl(H_0 \cup H_1 \cup H_2), |H_0 \cup H_1 \cup H_2|)$ must be added to $\mathcal{C}$ by step $|H_0 \cup H_1 \cup H_2|$ of the algorithm.

This process can be repeated (by replacing Z_i with Z_{i+1} and $H_0 \cup \ldots \cup H_i$ with $H_0 \cup \ldots \cup H_{i+1}$ in the previous paragraph) until $cl(H_0 \cup H_1 \cup \ldots \cup H_t) = V$ for some $t \in [n]$. Since the algorithm terminates when $cl(H_0 \cup \ldots \cup H_t) = V$ and returns $|H_0 \cup \ldots \cup H_t|$, and since $Z(G) = Z(G; H_0) = Z(G; H_0 \cup H_1) = \ldots = Z(G; H_0 \cup \ldots \cup H_t) = |H_0 \cup \ldots \cup H_t|$, it follows that the Wavefront Algorithm returns $Z(G)$. □

Theorem 2. *At any step s of the Wavefront algorithm,*

$$|\mathcal{C}| \leq \sum_{i=1}^{s} \frac{n!}{(n-i)!i!}, \tag{1}$$

and this bound is tight.

Proof. Multiple sets that have the same closure are not added to $\mathcal{C}$. Since all permutations of a set have the same closure, at most $\binom{n}{s}$ new sets can be added to $\mathcal{C}$ in step s. Thus, after step s, $|\mathcal{C}|$ is bounded as in (1). The worst case performance of Wavefront is realized in empty graphs (after n steps), since the closure of any set of vertices is the set itself, and since every combination of vertices must be checked before a zero forcing set is found. □

As shown in Theorem 2, in the worst case, the Wavefront algorithm is no better than enumerating all possible subsets of vertices; graphs in which very few vertices can be forced by sets with fewer than $Z(G)$ elements (e.g. stars) also lead to poor performance. However, Wavefront performs much better than the brute force algorithm when the closures of subsets of vertices are larger than the original subsets. This improvement comes from the fact that when some vertices being forced have no uncolored neighbors, they are no longer possible choices to add to the sets in $\mathcal{C}$.

While Wavefront could potentially be modified to create connected forcing sets, such a modification would eliminate the computational advantages of the algorithm. Wavefront has good performance because it only stores optimal forcing sets for certain subgraphs of G, and then builds the optimal forcing sets of larger subgraphs by adding neighborhoods of vertices containing an uncolored vertex. However, an optimal forcing set for a subgraph may not be connected to other vertices that must be added in order to force the entire graph. Thus, to be useful for finding connected forcing sets, Wavefront would have to store more than just the optimal forcing set for each subgraph, and its performance would suffer as a result. For these reasons, Wavefront will not be a viable method for solving the connected zero forcing problem without significant alterations that are beyond the scope of this paper.

3.2 Branch-and-Bound Algorithm

We now propose a combinatorial algorithm based on branch-and-bound for connected zero forcing. The central part of this method is the generation of connected subgraphs, which is achieved through a variant of reverse search. It was shown by Avis and Fukuda [5] that the reverse search technique generates all connected induced subgraphs of a graph $G = (V, E)$. In particular, this is done by starting from a subset containing each single vertex of G, and recursively adding to that subset one of its neighboring vertices. This process defines a tree T, whose leaves are the connected induced subgraphs of G. More precisely, given a subset $S \subset V$, the choice of whether or not a certain neighbor of S is in the connected subgraph gives two branches of a subtree of T descending from a node representing S. T can be traversed, e.g. by depth-first search, to generate all connected subgraphs of G; in addition, each subtree of T only includes subsets of vertices that are larger than the subset represented by the root of the subtree. Thus, once a connected zero forcing set of a certain size is found, all branches of the tree that lead to subsets of equal or greater size can be pruned. This idea is described formally in Algorithm 2.

Algorithm 2: Branch-and-bound algorithm for Connected Zero Forcing

```
Data: G = (V, E); three sets of vertices R, S, N(S); a constant ℓ
Result: A minimum connected forcing set of G
function GenerateSubgraph
    if R, S, and N(S) are not initialized then
        R = V, S = ∅, N(S) = ∅, ℓ = |V|;
    if S = ∅ then C = R;
    else C = R ∩ N(S);
    if C = ∅ then
        if cl(S) = V then
            ℓ = |S|;
            return S;
        else return;
    else
        Choose any v ∈ C;
        GenerateSubgraph(R\{v}, S, N(S), ℓ);
        if |S| < ℓ − 1 then GenerateSubgraph(R\{v}, S ∪ {v}, N(S) ∪ N(v), ℓ);
```

The correctness of Algorithm 2 follows from the fact that all connected induced subgraphs whose order is less than or equal to the cardinality of an already-discovered connected zero forcing set are enumerated. Since a graph G may have exponentially-many connected induced subgraphs with fewer than $Z_c(G)$ vertices, in the worst case this algorithm is no better than brute force. However, as with Wavefront, Algorithm 2 performs much better in practice, when G has relatively few connected induced subgraphs or when large parts of the search tree are pruned.

4 Integer Programming Approaches

In this section, we describe several integer programming formulations and solution strategies for computing the zero forcing and connected forcing numbers of a graph, and for finding the smallest set which forces a graph within a specified number of timesteps. The presented formulations come from two distinct perspectives on zero forcing. The first perspective is a straightforward model of zero forcing as a dynamic graph infection process. This approach incorporates the dynamic nature of the forcing process by using the vertices forced at each timestep to determine the vertices that can be forced in the next timestep. The second perspective uses the theory of zero forcing forts introduced by Fast and Hicks [30], and models zero forcing as a type of set-covering problem which does not depend on timesteps.

4.1 Infection Perspective

In the formulation of zero forcing as a dynamic process, each edge of the given graph $G = (V, E)$ is replaced by two directed edges with opposite directions. A binary variable s_v indicates whether vertex v is in the forcing set; an integer variable x_v with values between 0 and T indicates at which timestep vertex v is forced, where T is the maximum difference between the forcing times of two vertices; finally, a binary variable y_e for each directed edge $e = (u, v)$ indicates whether u forces v. For a directed edge e and vertex v, we use the notation $e \to v$ to indicate that v is the head of e.

Model 1 *IP model for Zero Forcing based on infection*

$$\begin{aligned}
\min\ & \sum_{v \in V} s_v \\
\text{s.t.:}\ & s_v + \sum_{e \to v} y_e = 1 \qquad \forall v \in V && (2)\\
& x_u - x_v + (T+1)y_e \leq T \qquad \forall e = (u, v) \in E && (3)\\
& x_w - x_v + (T+1)y_e \leq T \qquad \forall e = (u, v) \in E, \forall w \in N(u) \backslash \{v\} && (4)\\
& x \in \{0, \ldots, T\} \\
& s, y \in \{0, 1\}
\end{aligned}$$

Theorem 3. *The optimum of Model 1 is equal to $Z(G)$.*

Proof. Let Z be a zero forcing set, fix some chronological list of forces, and let $\mathcal{F}$ be the associated set of forcing chains. Since Z is a zero forcing set, each vertex v of G is either in Z (i.e. $s_v = 1$) or is forced by some other vertex of G (i.e. $y_e = 1$ and v is the head of e). Thus, constraint (2) must be satisfied. Now, let x_v be the timestep in which v is forced. Since a vertex cannot force until all-but-one of its neighbors are forced, it follows that for every edge $e = (v, w)$ for which $y_e = 1$, v must be forced before w and thus $x_v < x_w$. Likewise, $x_i < x_w$ for all neighbors i of v. Thus, constraints (3) and (4) are satisfied. If $y_e = 0$, then constraints (3) and (4) are satisfied since T is the maximum difference between the forcing times of two vertices. Thus, the constraints are valid for any zero forcing set and associated set of forcing chains.

Conversely, let (s, x, y) be a feasible solution of Model 1 for the given graph G, and let Z be the set of all vertices for which $s_v = 1$. Let $\mathcal{F}$ be the set of paths induced by the edges for which $y_e = 1$. For any edge $e = (v, w)$ for which $y_e = 1$, by constraints (3) and (4), it must be the case that v can force w through some number of applications of the color change rule. Since we also have that every vertex is either in Z or has an incoming edge with $y_e = 1$, it follows that every vertex in G is eventually forced. Therefore, Z is a zero forcing set of G and $\mathcal{F}$ is a set of forcing chains associated with Z. □

Note that for any graph, since at least one vertex is forced at each timestep, it follows that $T < n$; it is possible to further bound T if the graph is assumed to have certain properties (see, e.g., [16,22,30,36,46] for results on the propagation time T).

The main advantage of Model 1 is that the numbers of constraints and variables in this formulation are polynomial in n; therefore, the integer program can be solved directly, without row and column generation. Another useful feature of this model is that it not only finds the zero forcing number and a minimum zero forcing set of G, but it also gives a set of forcing chains associated with the forcing set. However, the downfall of Model 1 is its reliance on constraints (3) and (4), which are of big-M form; in general, these constraints lead to poor performance of this model (as shown in Section 5).

Despite its ineffectiveness at computing $Z(G)$, Model 1 can be used for a different purpose – namely, for finding a zero forcing set which forces G within a specified number of timesteps, and has minimum cardinality subject to that property. This result can be achieved by fixing T in Model 1 to be the maximum acceptable propagation time. As mentioned above, throttling the propagation time has been previously investigated (cf. [16,22,30,36,46]) from a combinatorial standpoint, but to our knowledge, this is the first computational tool for this problem for general graphs. Our experiments show that when T is fixed to be a smaller constant than n, Model 1 solves faster and is able to handle larger graphs.

4.2 Fort Covering Perspective

Our next formulation models zero forcing as a set-covering problem which does not depend on timesteps and does not rely on big-M constraints. In Model 2, the binary variable s_v again indicates whether vertex v is in the zero forcing set; $\mathcal{B}$ is the set of all forts in the given graph $G = (V, E)$.

Model 2 *IP model for Zero Forcing based on forts*

$$\begin{aligned} \min \ & \sum_{v \in V} s_v \\ \text{s.t.:} \ & \sum_{v \in B} s_v \geq 1 \qquad \forall B \in \mathcal{B} \\ & s \in \{0,1\} \end{aligned} \tag{5}$$

Theorem 4. *The optimum of Model 2 is equal to $Z(G)$.*

Proof. Let s be an optimal solution of Model 2, and let Z be the set of all vertices v for which $s_v = 1$. Suppose for contradiction that Z is not a zero forcing set of G, which means $cl(Z) \neq V$. If any vertex $u \in cl(Z)$ was adjacent to exactly one

vertex $v \in V \backslash cl(Z)$, then u could force v, contradicting the definition of $cl(Z)$. Thus, $V \backslash cl(Z)$ is a fort, and it does not contain any vertex of Z; this means constraint (5) is violated. It follows that Z must be a zero forcing set of G.

Conversely, let Z be a minimum zero forcing set of G. Suppose for contradiction that there exists a fort F which does not contain any element of Z. In order for the first vertex v of F to be forced, at some timestep, v must be the only neighbor of some colored vertex outside F. However, since F is a fort, any vertex outside F which is adjacent to v is also adjacent to another (uncolored) vertex in F. Thus, v cannot be forced, which contradicts Z being a forcing set. It follows that every fort contains an element of Z, so Z is a feasible solution of Model 2. □

In contrast to Model 1, Model 2 has the advantage of not having big-M constraints. However, the main issue with Model 2 is that since a graph could have an exponential number of forts, the solution methodologies for Model 2 must use constraint generation (see [23] for an introduction to this approach). The usefulness of constraint generation depends on the development of a practical method for finding violated constraints. One way to generate violated constraints is to simply find the closures of solutions to Model 2; if the closure of a solution is not the entire graph, then the vertices outside the closure form a fort, which gives a violated constraint.

Another method to find violated forts is to use the auxiliary integer program given in Model 3. For this model, we define the set S to be the set of all vertices for which $s_v = 1$ in the current optimal solution of Model 2. Note that since the value for each s_v is taken from the current optimal solution of Model 2, S is constant for Model 3. The x_v variables indicate whether vertex v is in the fort. The constant vector c enforces some desired property; for example, in our experiments, we set $c_v = 0.0001$ for each v in order to make the model find minimum size forts.

Model 3 *IP model for finding forts*

$$\begin{aligned}
\min\ & \sum_{v \in V} c_v x_v \\
\text{s.t.:}\ & \sum_{v \in V} x_v \geq 1 && (6)\\
& x_w - x_v + \sum_{a \in N(w) \backslash v} x_a \geq 0 \qquad \forall (v,w) \text{ with } v \in V, w \in N(v) && (7)\\
& x_v = 0 \qquad \forall v \in cl(S) && (8)\\
& x \in \{0,1\}
\end{aligned}$$

Theorem 5. *Model 3 finds a minimum weight violated fort with respect to the weights given by c.*

Proof. Let B be a violated fort of G; let $x_v = 1$ for $v \in B$, and $x_v = 0$ for $v \notin B$. Since a fort is non-empty by definition, B must contain at least one vertex; therefore, constraint (6) of the model is satisfied. Again by the definition of a fort, any neighbor of a vertex in B must either be in B or have at least one other neighbor in B; therefore, constraint (7) is satisfied. Since B is a violated fort, no vertex in B can be in $cl(S)$; therefore, constraint (8) is satisfied.

Conversely, let x be a solution of Model 3, and let B be the set of vertices of G for which $x_v = 1$. By constraint (6), B is not empty. By constraint (7), every neighbor of a vertex in B must either be in B, or have at least two neighbors in B. Thus, B is a fort of G. Furthermore, by constraint (8), no vertex in B is in $cl(S)$; therefore B is a violated fort with respect to the solution S. □

Model 3 separates violated constraints for Model 2. Theoretically, it is still difficult to solve; however, there is precedent in literature for using an integer programming separation method (see, e.g., [4,32]). In our computational experiments, Model 3 solved relatively quickly, and the forts found using this method are smaller and more effective at solving Model 2 than those found by the closure method described previously.

Since Model 2 is a set-covering problem, we can use the theory of Balas and Ng [6] on the set-covering polytope to explain why the forts generated by Model 3 are more effective than those found by the closure method. In particular, we restate a theorem from [6] in terms of forts; this result gives necessary and sufficient conditions for inequalities with a right-hand-side of one to be facet-inducing.

Theorem 6. [6] *Given a fort B, the inequality $\sum_{v \in B} s_v \geq 1$ defines a facet of the zero forcing polytope if and only if the following two conditions hold:*

$$\nexists A \subset B \text{ s.t. } A \text{ is a fort,} \tag{9}$$

$$\forall v \notin B, \exists w \in B \text{ s.t. if } A \subset B \cup \{v\} \text{ is a fort and } v \in A, \text{ then } w \in A. \tag{10}$$

Condition (9) explains why Model 3 performs better than the closure method. The latter makes no effort to minimize the size of the generated forts; thus, they are unlikely to satisfy (9) and be facet-inducing. On the other hand, Model 3 finds minimum size violated forts, which satisfy (9). However, condition (10) is not necessarily satisfied by either method.

This motivates further investigation of ways to add facet-inducing inequalities to Model 2, which we address next. Note that if (10) is violated for a fort F, then there exist p forts $A_1, \ldots, A_p \subset F \cup \{v\}$ such that there is a vertex v with $v \in A_i$, $1 \leq i \leq p$, but $\bigcap_{i \in [p]} (A_i \backslash \{v\}) = \emptyset$. Observe also that the fort constraints given by F and A_i for $1 \leq i \leq p$ can be combined to give the valid cut $\sum_{i \in F \cup \{v\}} x_i \geq 2$. This valid cut is found by first summing the fort constraints corresponding to F and all the A_i; since $\bigcap_{i \in [p]} A_i = \emptyset$, the coefficients of each vertex variable in the sum is at most p, but the right-hand-side is $p + 1$. Thus, the valid cut can be obtained by dividing through by p and taking the ceiling of each coefficient. Moreover, the magnitude of p is bounded as follows.

Theorem 7. *Suppose there exist p forts $A_1, \ldots, A_p \subset F \cup \{v\}$ such that there is a vertex v with $v \in A_i$ for $1 \leq i \leq p$, but $\bigcap_{i\in[p]}(A_i \backslash \{v\}) = \emptyset$. Then p can be chosen to be at most $|F|$.*

Proof. Suppose there exist $q > |F|$ forts that satisfy the required properties. Then, for each vertex $w \in F$, choose one fort among $A_1, \ldots, A_q$ that does not contain w; since $\bigcap_{i\in[q]}(A_i \backslash \{v\}) = \emptyset$, such a fort must exist for each w. This collection consists of at most $|F|$ forts, whose intersection is empty except for v. Thus, it is possible to choose $p \leq |F|$ forts which satisfy the required properties. □

Given the above theory, Model 4 can be used to check whether a fort generated by Model 3 is facet-inducing. If the generated fort is not facet-inducing, then the valid cut generated as described above can be added instead of a fort constraint. In Model 4, the variable x_{ij} indicates whether vertex j is chosen to be in fort i, and the variable y_i indicates whether fort i is empty.

Model 4 *IP model for checking if a fort is facet-inducing*

$$\min \sum_{i\in[|F|]} y_i$$

$$\text{s.t.:} \sum_{v\in V\backslash F} x_v = 1 \tag{11}$$

$$\sum_{v\in V\backslash F} x_{iv} = y_i \qquad \forall i \in [|F|] \tag{12}$$

$$x_{iv} \leq x_v \qquad \forall i \in [|F|], \forall v \in V\backslash F \tag{13}$$

$$\sum_{i\in[|F|]} x_{iw} \leq \sum_{i\in[|F|]} y_i - 1 \qquad \forall w \in F \tag{14}$$

$$x_{iw} - x_{iu} + \sum_{a\in N(w)\backslash\{u\}} x_{ia} \geq 0 \quad u \in V, w \in N(v), \forall i \in [|F|] \tag{15}$$

$$x_{iw} \leq y_i \qquad \forall i \in [|F|], \forall w \in V \tag{16}$$

$$x, y \in \{0, 1\}$$

Theorem 8. *If Model 4 is infeasible, and F is a minimum size fort, then the fort F is facet-inducing. If Model 4 has an optimal solution, then the set of forts with $y_i = 1$ shows that F is not facet-inducing by condition (10) of Theorem 6.*

Proof. Suppose a minimum size fort F is not facet-inducing. Since F has minimum size, it must satisfy condition (9), so F can only violate (10). Therefore, there must exist p forts $A_1, \ldots, A_p \subset F \cup \{v\}$ such that there is a vertex v with $v \in A_i$ for $1 \leq i \leq p$ but $\bigcap_{i\in[p]}(A_i \backslash v) = \emptyset$. By Theorem 7, we can assume that $p \leq |F|$. Let $y_i = 1$ for $1 \leq i \leq p$, and let $y_i = 0$ otherwise; let $x_{iw} = 1$ if w

is in fort A_i for $1 \leq i \leq p$ and $x_{iw} = 0$ otherwise; let $x_v = 1$ and $x_{iv} = 1$ if $1 \leq i \leq p$. All other variables are set to 0. Now, observe that (11) is satisfied because $x_w = 0$ for all $w \neq v$, and $x_v = 1$. Constraint (12) is satisfied because each fort A_i contained v, and (13) is satisfied because x_{iv} is either 0 or 1 and $x_v = 1$. Constraint (14) is satisfied because $\bigcap_{i\in[p]}(A_i \backslash v) = \emptyset$. Constraint (15) is satisfied because each A_i was a fort, and (16) is satisfied because the x variables are chosen to be 1 only for the forts with y variables chosen to be 1. Hence, (x, y) is a feasible solution to Model 4. Thus, if Model 4 is infeasible, then the fort F must be facet-inducing.

Conversely, if Model 4 has an optimal solution, then defining $A_i = \{w \in V : x_{iw} = 1\}$ gives a set of forts which shows that F does not satisfy condition (10) of Theorem 6, and hence F is not facet-inducing. □

Instead of determining with certainty whether a fort constraint is facet-inducing using Model 4, we can also determine this characteristic in a heuristic manner. In particular, Model 4 can be simplified by limiting the number of forts that can be chosen, i.e., requiring at most 2 forts instead of $|F|$ forts. This simplification is given in Model 5 below.

Model 5 *IP model for checking if a fort is "likely" to be facet-inducing*

$$\min \sum_{i\in[|F|]} y_i$$

$$\text{s.t.:} \sum_{v\in V\backslash F} x_v = 1 \tag{17}$$

$$\sum_{v\in V\backslash F} x_{iv} = 1 \qquad \forall i \in \{1, 2\} \tag{18}$$

$$x_{1v} \leq x_v \qquad \forall v \in V\backslash F \tag{19}$$

$$x_{1v} + x_{2v} \leq 1 \qquad \forall v \in F \tag{20}$$

$$x_{iw} - x_{iu} + \sum_{a\in N(w)\backslash\{u\}} x_{ia} \geq 0 \qquad u \in V, w \in N(v), \forall i \in \{1, 2\} \tag{21}$$

$$x \in \{0, 1\}$$

Model 5 is easier to solve than Model 4, but it does not determine with certainty whether a fort constraint is facet-inducing. If a feasible solution is found for Model 5, then the fort constraint is definitely not facet-inducing; if the model is infeasible, then the constraint may or may not be facet-inducing. Nevertheless, our computational experiments show that Model 2 generally has better performance when combined with Model 5 than when combined with Model 4.

4.3 Fort Covering Extended

Another way to improve the formulation of Model 2 comes from the observation that any zero forcing set must contain some vertex together with all-but-one of its neighbors. This idea is expressed in Model 6, with the addition of binary variables z_v, which indicate that vertex v and all-but-one of its neighbors belong to a zero forcing set. Hence, the z_v variables have a cost of $|N(v)|$ in the objective function, and at least one of them is required to be positive (enforced by constraint (23)). Note that if a vertex w is in the closure of $N[v]$ and z_v is positive, then w will be forced in the corresponding solution. Therefore, s_w and z_v will never both be positive; this is enforced by constraint (24). Finally, the fort constraints (constraint (5) of Model 2) are modified to (22) to allow satisfaction by z_v variables. Despite the increased number of variables, Model 6 performs better in our experiments than Model 2.

Model 6 *IP model for Zero Forcing based on forts with neighborhood variables*

$$\begin{aligned}
\min\ & \sum_{v\in V} |N(v)| z_v + \sum_{v\in V} s_v \\
\text{s.t.:}\ & \sum_{v\in B} \Big(s_v + \sum_{v\in cl(N[w])} z_w\Big) \geq 1 \qquad \forall B \in \mathcal{B} && (22)\\
& \sum_{v\in V} z_v \geq 1 && (23)\\
& s_w + z_v \leq 1 \qquad \forall v \in V, \forall w \in cl(N[v]) && (24)\\
& s, z \in \{0,1\}
\end{aligned}$$

Given the additional variables in Model 6, the constraint generation models must also be expanded to generate violated forts. Instead of minimizing the number of vertices in the fort as in Model 3, our experiments showed better performance when we minimize the number of vertices in the fort that are adjacent to vertices outside of the fort. Such *minimum border forts* can be found using the integer program in Model 7. In this model, c_v is a penalty coefficient to penalize vertices that are adjacent to the closure, and b_v is a binary variable that indicates whether the vertex v is adjacent to vertices outside of the fort. S is the set containing every vertex v such that either $s_v = 1$ in the current solution, or v is in $N[w]$ for some w with $z_w = 1$ in the current solution. Constraint (27) ensures that the b_v variables correctly indicate whether v is on the border of the fort.

4.4 Fort Covering for Connected Zero Forcing

In this section, we adapt the integer programming models introduced previously to the connected forcing problem, by adding constraints to enforce connectivity on the chosen zero forcing set. We focus on adding connectivity constraints to

Model 7 *IP model for finding minimum border forts*

$$\min \sum_{v \in V} c_v b_v$$

$$\text{s.t.:} \sum_{v \in V} x_v \geq 1 \tag{25}$$

$$x_w - x_v + \sum_{a \in N(w) \setminus v} x_a \geq 0 \qquad \forall v, w \text{ with } v \in V, w \in N(v) \tag{26}$$

$$|N(v)|x_v - |N(v)|b_v - \sum_{a \in N(v)} x_a \leq 0 \qquad \forall v \in V \tag{27}$$

$$x_v = 0 \qquad \forall v \in cl(S) \tag{28}$$

$$x, b \in \{0, 1\}$$

Model 2, because it is the best performing model from the previous sections that allows us to ensure connectivity. The neighborhood variables in Model 6 make it difficult to enforce connectivity since one vertex in the neighborhood is not in the zero forcing set.

Drawing from the literature on connected dominating sets and connected power dominating sets, there are multiple ways of enforcing connectivity in integer programs. Fan and Watson [29] compared Miller-Tucker-Zemlin (MTZ) constraints, Martin constraints, single-commodity flow constraints, and multi-commodity flow constraints, and found that the MTZ constraints provide the best computational performance for both the connected dominating set and connected power dominating set problems. Another method of enforcing connectivity is to add a,b-separation cutting planes when needed, in order to cut off disconnected solutions. This method has been used for connected dominating sets [14], Steiner trees [31], and forest planning problems [21] (see also [45] for conditions that cause such inequalities to induce facets of the connected subgraph polytope). In view of these results, we explore the effectiveness of MTZ constraints and a,b-separation inequalities for enforcing connectivity in the connected zero forcing problem.

MTZ constraints were originally introduced by Miller, Tucker, and Zemlin [41] in relation to the Traveling Salesman Problem; see also [29,42] for further work on these constraints, which we adapt in our implementation. The basic idea of MTZ constraints is to enforce the existence of a directed spanning tree in the subgraph induced by the chosen vertices. Our implementation follows Fan and Watson's [29] explanation of the method introduced by Quintão, da Cunha, Mateus, and Lucena [42]. Two new vertices labeled α and β are added to the given graph $G = (V, E)$, along with a set E_{new} of edges containing a directed edge from each of the two new vertices to all the original vertices; E_{new} also contains a directed edge from α to β. In the modified graph, the vertices which are not chosen to be in a forcing set will have a positive edge variable coming into them from α, while β will have a positive edge variable going to the root of the directed spanning tree of the chosen connected zero forcing set.

Model 8 combines Model 2 with the MTZ constraints. In Model 8, (29) is the original fort cover constraint from Model 2; the rest of the constraints are the MTZ constraints. In particular, (30) ensures that there is an edge chosen from β to some vertex that will be the root of the directed spanning tree of the zero forcing set; (31) ensures that each vertex has an incoming edge. Constraint (32) ensures that vertices connected to α cannot be used to connect to any other vertices; (33) and (34) ensure that there are no cycles in the chosen edges. Constraint (35) ensures that vertices chosen to be in the forcing set must be in the spanning tree instead of connected to α. The rest of the constraints impose bounds on the variables.

Model 8 *IP model for Connected Zero Forcing using MTZ constraints*

$$\min \sum_{v \in V} s_v$$

$$\text{s.t.:} \sum_{v \in B} s_v \geq 1 \qquad \forall B \in \mathcal{B} \tag{29}$$

$$\sum_{v \in V} y_{\beta,v} = 1 \tag{30}$$

$$\sum_{i:(i,v) \in E} y_{i,v} = 1 \qquad \forall v \in V \tag{31}$$

$$y_{\alpha,v} + y_{v,i} \leq 1 \qquad \forall (v,i) \in E \tag{32}$$

$$(n+1)y_{i,v} + u_i - u_v + (n-1)y_{v,i} \leq n \qquad \forall (v,i) \in E \tag{33}$$

$$(n+1)y_{i,v} + u_i - u_v \leq n \qquad \forall (v,i) \in E_{new} \tag{34}$$

$$x_v = 1 - y_{\alpha,v} \qquad \forall v \in V \tag{35}$$

$$y_{\alpha,\beta} = 1$$

$$u_\alpha = 0$$

$$1 \leq u_v \leq n+1 \qquad \forall v \in V \cup \{\beta\}$$

$$s, y \in \{0,1\}$$

$$u \in \mathbb{Z}$$

A solution of Model 8 is a minimum connected forcing set by Theorem 4 and by the fact that MTZ constraints impose connectivity of the selected set of vertices. Although the number of variables in Model 8 is more than triple that of Model 2, and some of the variables are integer instead of binary, we found that Model 8 performs very well in practice.

The second method for enforcing connectivity that we explore relies on adding valid inequalities which cut off disconnected solutions, rather than on using additional variables. These valid inequalities added to Model 2 are known as a,b-separation inequalities; the idea behind them is that if a set $C \subset V$ is a vertex cut separating vertices a and b in a graph $G = (V, E)$, and both a and b are chosen to be in connected zero forcing set, then some vertex from C must also

belong to this forcing set. Model 9 gives the complete integer programming formulation for connected zero forcing using a,b-separation inequalities. Constraint (36) is the original fort cover constraint from Model 2, and (37) expresses the a,b-separation inequalities.

Model 9 *IP model for Connected Zero Forcing using a,b-separator constraints*

$$\begin{aligned}
\min\ & \sum_{v\in V} s_v \\
\text{s.t.:}\ & \sum_{v\in B} s_v \geq 1 \qquad \forall B \in \mathcal{B} && (36)\\
& s_a + s_b - \sum_{v\in C} s_v \leq 1 \qquad \forall a,b \in V \text{ with } C \text{ an } a,b\text{-separator} && (37)\\
& s \in \{0,1\}
\end{aligned}$$

The a,b-separation inequalities can be separated efficiently using the observation that if the chosen zero forcing set Z is not connected, then the set $C = V\backslash Z$ must be a vertex cut separating at least two vertices $a \in Z$ and $b \in Z$. However, as was pointed out by Buchanan et al. [14], the resulting vertex cuts in such an implementation are likely larger than necessary. Since the decision variable for each vertex in C appears in these constraints, the constraints are stronger when the size of the vertex cut S is minimized. Therefore, vertices could be deleted from a vertex cut of G until it becomes inclusion-minimal. For the dominating set problem (which was the focus of [14]), a valid cutting plane can be obtained from a vertex cut; however, a zero forcing set does not have to be dominating. Therefore, we also require that the vertex cut must be an a,b-separator for $a, b \in Z$. In Algorithm 3, we give a modified version of the approach of Buchanan et al. [14], which ensures the resulting vertex cut will be an a,b-separator.

Algorithm 3: a,b-Separator Algorithm

Data: A graph $G = (V, E)$, two vertices $a, b \in V$, and an a,b-separator $C \subset V$
Result: An inclusion-minimal a,b-separator $C' \subset C$
$C' = \{v \in C : \exists w \notin C, \{v, w\} \in E\}$;
$\mathcal{S} = \{S : S$ is a connected component of $G[V\backslash C']\}$;
for $v \in C'$ **do**
  if *v is not adjacent to $S_a, S_b \in \mathcal{S}$ with $a \in S_a$ and $b \in S_b$* **then**
    $C' = C'\backslash\{v\}$;
    $\mathcal{S}^v = \{S \in \mathcal{S}$: v is adjacent to $S\}$;
    $S_{new} = \bigcup_{S\in\mathcal{S}^v} S \cup \{v\}$;
    $\mathcal{S} = S_{new} \cup (\mathcal{S}\backslash\mathcal{S}^v)$;

5 Computational Results

This section presents implementation details and computational results for finding minimum zero forcing sets and minimum connected forcing sets using the methods described thus far. In particular, for the zero forcing problem, we compare

- the Wavefront algorithm (Algorithm 1),
- the Infection model (Model 1),
- the Fort Cover model without generation of facet-inducing forts (Model 2 together with Model 3),
- the Fort Cover model with simplified generation of facet-inducing forts (Model 2 together with Model 3 and Model 5),
- the Extended Fort Cover model (Model 6 together with Model 7).

For the connected forcing problem, we compare

- the Brute Force algorithm,
- the Branch-and-Bound algorithm (Algorithm 2),
- the Fort Cover model with MTZ constraints and fort generation (Model 8 together with Model 3),
- the Fort Cover model with a,b-separator constraints and fort generation (Model 9 together with Algorithm 3 and Model 3).

5.1 Implementation Details

Our computational results were obtained on a Dell Precision T1650 workstation with a 3.3 GHz Intel Core i3-2120 CPU, 3.7 GB of RAM, and Red Hat Enterprise Linux version 6.6. Integer programs were solved using Gurobi version 5.5.0 set to use a single thread. The Brute Force, Branch-and-Bound, and Wavefront algorithms were implemented in C++ and compiled with g++ version 4.8. Implementations of all programs and models used can be obtained at `https://github.com/calebfast/zero_forcing`.

We tested the different solution approaches on three classes of random graphs: cubic graphs, WS$(5, 0.3)$ graphs, and WS$(10, 0.3)$ graphs. We used our own C++ implementation to generate random cubic graphs, and we used the connected Watts-Strogatz graph generator from the NetworkX version 1.8.1 package in Python 2.7.6 to generate WS$(5, 0.3)$ and WS$(10, 0.3)$ graphs. For each family of graphs, we generated five random instances with $10i$ vertices for incrementing values of i, and tested each algorithm starting from the smallest graphs until all five graphs of a certain size could not be solved within 2 hours. We did not generate WS$(10, 0.3)$ graphs with 10 vertices, since those are just complete graphs.

When solving Model 2 without generation of facet-inducing forts, a maximal set of disjoint minimum size forts was added to the formulation using Model 3 before solving. Other violated fort constraints were added to the model using a MIPSOL callback, which is invoked by Gurobi whenever it finds a new integral

incumbent solution. When a violated fort is found, the MIPSOL callback adds that fort to the formulation as a lazy constraint (to enable lazy constraints, the "PreCrush" and "LazyConstraints" parameters of Gurobi were both set to 1). If no more violated forts are found, then Gurobi terminates with an optimal solution.

Similarly, when solving Model 2 with generation of facet-inducing forts, if a generated fort F is not facet-inducing, then the valid cut associated with the forts that show F is not facet-inducing is added instead of F. Our computational results in Table 1 show that while checking for facet-inducing forts provides a small benefit in average running time and reduces the number of forts that must be generated, it is not effective enough to increase the size of the instances that can be solved within 2 hours. Because checking for facet-inducing forts provided no consistent benefit, subsequent models using fort constraints do not check whether the generated forts are facet-inducing.

Model 6 was solved similarly to Model 2. Violated minimum border forts were generated using Model 7, and a maximal set of disjoint forts was added to the formulation before solving. In addition, for each $v \in V$, the fort $V \backslash cl(N[v])$ was also added to the formulation before solving. In our implementation, all values of c_v in Model 7 were set to 0.0001.

Model 8 was solved exactly like Model 2, with the addition of the MTZ constraints. Violated forts are added to the formulation by the MIPSOL callback as lazy constraints; when no more violated forts are found, Gurobi terminates with an optimal solution.

In Model 9, both fort constraints and a,b-separation inequalities are added to the model using a MIPSOL callback. This callback first generates a minimum size violated fort by solving Model 3, and adds it to the formulation as a lazy constraint. When no more violated forts are found, the callback checks whether the current solution is connected. If it is not connected, Algorithm 3 finds a minimal separator contained in the separator given by the vertices outside the current solution; the corresponding a,b-separation inequality is then added to Model 9 as a lazy constraint. When no violated forts are found and the solution is connected, Gurobi terminates with an optimal solution.

For all these models, the parameters not mentioned in the discussion above were left to their defaults in Gurobi. Some testing showed that tuning certain parameters (such as the branching direction (BranchDir), aggressiveness of cut generation (Cuts), or the focus of the solver (MIPFocus)) could improve performance on some specific instances, but not in general. The branching strategy was also left to the Gurobi default.

Table 1 gives the average runtimes of the different zero forcing algorithms for the graphs tested; Table 2 gives the average runtimes of the different connected forcing algorithms; Table 3 gives the average runtimes of Model 1 modified to run with a limited propagation time. The reported times reflect the time taken by Gurobi to optimize the relevant models; they do not include the time necessary for data input and setting up the Gurobi models.

Table 1. Average running times for zero forcing algorithms on different random graphs. All times are in seconds. Asterisks indicate that not all instances of the specified size were solved. In these cases, the reported result is the average time for the instances that were successfully solved. 'T' indicates that none of the instances were solved within 2 hours. Bold text indicates the algorithm with the best performance.

	$\|V\|$	$Z(G)$	Wavefront	Infection	FC w/facet	FC no facet	Ext. Cover
Cubic graphs	10	3.8	**0.0013**	0.0470	0.57	0.022	0.024
	20	5.2	**0.017**	77.43	0.46	0.32	0.11
	30	6.6	**0.18**	206.20*	1.88	1.73	0.64
	40	8.8	**2.79**	T	24.60	35.74	7.18
	50	9.2	**9.68**	T	200.56	274.04	40.33
	60	11.4	**227.02**	T	5986.44*	5925.79*	1813.21*
	70	12	**525.46**	T	T	T	T
	80	12	**681.13***	T	T	T	T
WS(5, 0.3) graphs	10	4.4	**0.0013**	0.70	0.197	0.048	0.047
	20	6.2	**0.018**	152.47	1.40	0.57	0.55
	30	7	**0.072**	5320.23*	5.79	4.75	3.58
	40	9.4	**1.32**	T	73.79	81.06	47.58
	50	10.8	**10.47**	T	2552.73*	2694.03	2234.27
	60	11.6	**69.30**	T	438.14*	475.11*	2387.36*
	70	14	**678.89**	T	T	T	T
	80	14.5	**1306.57***	T	T	T	T
WS(10, 0.3) graphs	20	12	**0.010**	T	100.48	10.47	11.52
	30	15.4	**0.11**	T	2492.30	1069.80	1370.70
	40	18	**0.93**	T	T	T	T
	50	21.8	**9.41**	T	T	T	T
	60	24.6	**65.51**	T	T	T	T
	70	27.4	**416.65**	T	T	T	T
	80	30	**2192.86***	T	T	T	T

5.2 Comparison and Discussion

For the zero forcing problem, the computational results in Table 1 show that the Wavefront algorithm performs best, followed by the Extended Fort Cover model and the standard Fort Cover model; all these approaches considerably outperform the Infection model. Moreover, the Fort Cover model which generates facet-inducing forts is similar in overall runtime to the Fort Cover model which does not generate facet-inducing forts. All three Fort Cover models can handle similarly sized graphs, although the Extended Fort Cover model solves the instances faster on average, especially in sparser graphs. The difference in results between cubic and Watts-Strogatz graphs indicates that the four integer programs are sensitive to the density and vertex degrees of the graphs, while the Wavefront algorithm is less sensitive to these changes, and is primarily affected by size of the graphs. While the Wavefront algorithm has the best performance

Table 2. Average running times for connected forcing algorithms on random graphs. All times are in seconds. Asterisks indicate that not all instances of the specified size were solved. In these cases, the reported result is the average time for the instances that were successfully solved. 'T' indicates that none of the instances were solved within 2 hours. Bold text indicates the algorithm with the best performance (measured first by the number of instances solved within two hours and then by the average runtime).

	$\|V\|$	$Z_c(G)$	Brute	B&B	MTZ	a,b-sep
Cubic graphs	10	3.8	0.007	**0.001**	0.02	0.01
	20	5.4	0.50	**0.0128**	0.12	0.09
	30	7.8	1334.21	**0.28**	0.76	0.49
	40	9.8	T	537.0	4.49	**3.28**
	50	10.4	T	21.97	**19.95**	79.79
	60	12	T	550.0	**56.60**	1141.84
	70	13.4	T	287.33*	**261.54**	3541.51
	80	15.6	T	2857.07*	**1869.82**	1685.30*
	90	15.2	T	3192.58*	**2952.18**	T
	100	16	T	T	**2566.91***	T
WS(5, 0.3)	10	4.4	0.011	**0.003**	0.07	0.05
	20	6.2	1.72	**0.21**	0.42	0.24
	30	7	90.62	**1.34**	2.73	2.15
	40	9.4	T	25.86*	25.47	**20.48**
	50	10.8	T	1102.20*	**236.47**	267.90
	60	11.8	T	2223.9*	**2199.54***	1876.90*
	70	13	T	T	**7180.12***	T
WS(10, 0.3)	20	12	112.63	16.49	26.21	**9.65**
	30	15.4	T	T	3768.18*	**850.94**
	40	18	T	T	T	T

for the zero forcing problem, it is less flexible than the integer programming methods to accommodating different types of constraints or related problems, such as assuring connectivity or limiting the propagation time.

For the connected forcing problem, the computational results show that the Branch-and-Bound algorithm performs best on small, sparse graphs, but is outperformed by the integer programming models as the size and density of the graphs increases. This is because the Branch-and-Bound algorithm relies on enumerating connected induced subgraphs, and larger, denser graphs have significantly more such subgraphs. The Fort Cover model with MTZ constraints solves the largest instances out of any method, including those for standard zero forcing. This result is somewhat surprising, since the a,b-separator constraints have outperformed the MTZ constraints on other problems such as connected domination [14]. For the densest graphs tested (the WS$(10, 0.3)$ graphs), the Fort Cover Model with a,b-separation constraints was faster than the model with

Table 3. Average running times for Model 1 with different propagation times. All times are in seconds. The ΔZ columns give the average difference between $Z(G)$ and the minimum forcing set for the given propagation time. Asterisks indicate that not all instances of the specified size were solved. In these cases, the reported result is the average time for the instances that were successfully solved. 'T' indicates that none of the instances were solved within 2 hours.

		$T=1$		$T=2$		$T=4$		$T=8$	
	$\lvert V\rvert$	Time	ΔZ	Time	ΔZ	Time	ΔZ	Time	ΔZ
Cubic graphs	10	0.014	2.4	0.088	1	0.066	0	0.062	0
	20	0.077	6	7.73	3.2	25.26	0.6	67.28	0
	30	0.23	10.2	438.61	6.2	499.89*	1.6*	139.99*	0.3*
	40	0.32	13.8	T	T	T	T	T	T
	50	0.57	19	T	T	T	T	T	T
	60	2.23	21.8	T	T	T	T	T	T
	70	4.88	26.8	T	T	T	T	T	T
	80	8.41	32*	T	T	T	T	T	T
	90	11.41	N/A	T	T	T	T	T	T
	100	24.70	N/A	T	T	T	T	T	T
WS$(5, 0.3)$ graphs	10	0.014	1.8	0.19	0.8	0.63	0.2	0.75	0
	20	0.11	6.4	21.25	3.8	145.27	1.6	111.32	0
	30	0.41	11.6	2376.11	7.2	T	T	T	T
	40	0.76	14.8	T	T	T	T	T	T
	50	2.39	19.8	T	T	T	T	T	T
	60	3.52	25.4	T	T	T	T	T	T
	70	9.78	28	T	T	T	T	T	T
	80	18.87	N/A	T	T	T	T	T	T
	90	32.27	N/A	T	T	T	T	T	T
	100	162.18	N/A	T	T	T	T	T	T
WS$(10, 0.3)$ graphs	20	1.78	3	947.10	1.2	2097.47*	0*	T	T
	30	81.24	6.4	T	T	T	T	T	T
	40	451.37	11	T	T	T	T	T	T
	50	4744.08*	14.3*	T	T	T	T	T	T
	60	2909.11*	17.8*	T	T	T	T	T	T
	70	3361.71*	22.8*	T	T	T	T	T	T
	80	2810.39*	27*	T	T	T	T	T	T

MTZ constraints, although both were able to handle the same sized graphs. All three of the nontrivial approaches are faster and able to handle larger graphs than the Brute Force approach.

The helpfulness of MTZ constraints can be clearly seen when comparing the number of forts required to solve Model 2 vs. Model 8; the results in Table 4 show that the addition of MTZ constraints drastically decreases the number of forts that must be generated. Some additional insights can be gained by looking at

the average percentage of time spent to generate violated forts by each method, as shown in Table 5. The percentage of time spent by Model 8 to generate forts is also significantly less than the analogous model without connectivity constraints. The results of Table 5 also show that the Extended Cover model spends a lower percentage of its runtime generating forts than the Fort Cover model. While both models spend the majority of their time generating forts, this percentage decreases when the instances reach a size where the methods start to fail, and then increases again as the size of failure is surpassed. This behavior indicates that at the size of failure, Gurobi was not able to generate as many incumbent solutions and spent more time trying to prove optimality. This behavior makes sense because as the set of fort constraints in the model approaches a set of forts necessary to solve the zero forcing problem, then generating additional feasible solutions of lower cost will become more difficult, and the solver will spend more time proving optimality.

Table 4. Comparison of the average number of forts required by the Fort Cover model with MTZ constraints vs. the average number of forts required by the original Fort Cover model. The averages reported here are only over the instances that were solved by the methods within 2 hours.

	Cubic		WS$(5, 0.3)$		WS$(10, 0.3)$	
$\|V\|$	MTZ	FC	MTZ	FC	MTZ	FC
10	4.4	17.6	20.6	22.8		
20	11.4	67.2	36.4	102.8	887.2	1133.8
30	9.4	153.2	79.8	245.2	6754.8	14024.8
40	8.4	1834.2	168.8	2334.2	T	T
50	14.6	4569	388.6	15801	T	T
60	45.6	27082	697.5	1489	T	T
70	65	T	554	T	T	T
80	178	T	T	T	T	T
90	156.8	T	T	T	T	T

When comparing Models 8 and 9, we see that the MTZ constraints perform better for the cubic graphs, which are relatively sparse. The two models showed similar performance for the WS$(5, 0.3)$ graphs, and the a,b-separation constraints were better for the WS$(10, 0.3)$ graphs. This makes sense because as the average degree of the vertices increases, the likelihood that a chosen subset of vertices will induce a connected graph increases. Therefore, for the WS$(10, 0.3)$ graphs, which have high average degree, the a,b-separation inequalities were usually not necessary. In such cases, the a,b-separation inequalities are not added to the model, and it is solved as a basic zero forcing problem.

Finally, Model 1 runs faster and is able to handle larger graphs when it is coupled with different bounds on the propagation time T. For small values of T,

Table 5. Average percentage of time spent generating forts for Fort Cover (FC) and Extended Cover (EC) models for zero forcing, and Fort Cover with MTZ constraints and a,b-separation inequalities for connected forcing. Asterisks indicate that not all instances of the specified size were solved within 2 hours. In these cases, the reported number is the percentage of the 2 hours that was used to generate forts.

$\|V\|$	Cubic		WS$(5, 0.3)$		WS$(10, 0.3)$	
	FC	EC	FC	EC	FC	EC
10	75.9	72.6	88.8	85.9		
20	92.6	56.5	93.4	91.1	91.4	88.9
30	89.9	74.9	89.6	86.2	52.4	39.9
40	77.9	55.4	71.9	67.3	36.1*	32.6*
50	73.1	58.9	46.0*	48.1	T	T
60	20.8*	23.8*	93.8*	86.6*	T	T
70	57.0*	71.1*	82.7*	93.1*	T	T
	MTZ	a,b	MTZ	a,b	MTZ	a,b
10	21.9	29.6	43.6	57.6		
20	32.3	47.7	34.7	62.7	33.0	75.3
30	6.7	24.8	36.7	63.5	6.9	24.8
40	0.8	7.9	12.8	27.9	T	T
50	1.4	7.3	7.5	16.9	T	T
60	1.9	4.1	3.1*	13.5*	T	T
70	1.3	9.4	2.2*	14.3*	T	T
80	0.3	7.8*	T	T	T	T
90	0.6	7.1*	T	T	T	T

the model solved for all available graphs, although it was still somewhat impaired by the graph density. As seen in the ΔZ columns of Table 3, the sizes of the zero forcing sets with bounded propagation time approach the sizes of the minimum zero forcing sets as T grows. However, the increase in T also causes increased runtime due to the big-M constraints in the model.

6 Conclusion and Future Work

This paper introduced new methods for computing the zero forcing and connected forcing numbers of graphs. We presented combinatorial algorithms, as well as integer programming formulations based on an infection perspective and a set-covering perspective of zero forcing. We explored several solution strategies for these models, drawing from different areas of integer programming and polyhedral theory, and we compared their performance on random cubic and Watts-Strogatz graphs. Our computational experiments show that the Wavefront algorithm generally outperforms the integer programming models for zero forcing, while our approaches for connected forcing were faster, and able to handle larger graphs than Wavefront and the zero forcing models. We also presented

an integer program for finding a set of minimum cardinality which forces a graph within a specified number of timesteps; this model performed very well for small propagation times.

Some of the difficulty in solving the zero forcing problem is due to the symmetry of solutions; this symmetry arises from the fact that for each zero forcing set and any associated set of forcing chains, another zero forcing set of equal size can be obtained by choosing the terminals of the forcing chains [8]. Such sets of vertices are nearly indistinguishable in many of the algorithms and formulations, and this symmetry is harder to detect than simple isomporphisms in the graph. Any method for dealing with the symmetry of zero forcing has the potential to drastically improve the performance of the integer programs presented in this paper. Therefore, a direction for future work is to focus on breaking this symmetry. Note that this symmetry is somewhat less prevalent in connected forcing, since the set of terminals of forcing chains associated with a connected forcing set is not always connected; this may be part of the reason why the connected variants of the integer programs performed better.

Another difficulty in solving the problems is owed to the lack of monotonicity in the relationships between certain parameters. One instance of non-monotonicity is witnessed in the relationship between the size of a zero forcing set and its propagation time: a monotone change in one does not necessarily imply a monotone change in the other. As another example, fort containment is not monotone: a fort A may contain a proper subset B which is not a fort, and B may contain a proper subset C which is a fort. Therefore, it is not necessarily easy to find a minimal (nor a minimum) fort. Future work may be aimed at defining certain types of sets which have similar properties as forts but are closed under containment. Subsequently, it may be easier to find violated constraints for a formulation based on a cover by such sets.

Acknowledgements

This work is supported by the National Science Foundation, Grants No. 1450681, CMMI-1300477, and CMMI-1404864.

References

1. A. Aazami. Hardness results and approximation algorithms for some problems on graphs. PhD thesis, University of Waterloo, 2008.
2. AIM Special Work Group. Zero forcing sets and the minimum rank of graphs. *Linear Algebra and its Applications*, 428(7): 1628–1648, 2008.
3. D. Amos, Y. Caro, R. Davila, and R. Pepper. Upper bounds on the k-forcing number of a graph. *Discrete Applied Mathematics*, 181: 1–10, 2015.
4. P. Avella, M. Boccia, and I. Vasilyev. Computational experience with general cutting planes for the set covering problem. *Operations Research Letters*, 37: 16–20, 2009.
5. D. Avis and K. Fukuda. Reverse search for enumeration. *Discrete Applied Mathematics*, 65: 21–46, 1996.

6. E. Balas and S.M. Ng. On the set covering polytope: I. All the facets with coefficients in $\{0, 1, 2\}$. *Mathematical Programming*, 43: 57–69, 1989.
7. F. Barioli, W. Barrett, S.M. Fallat, T. Hall, L. Hogben, B. Shader, P. van den Driessche, and H. van der Holst. Parameters related to tree-width, zero forcing, and maximum nullity of a graph. *Journal of Graph Theory*, 72(2): 146–177, 2013.
8. F. Barioli, W. Barrett, S. Fallat, H.T. Hall, L. Hogben, B. Shader, P. van den Driessche, and H. van der Holst. Zero forcing parameters and minimum rank problems. *Linear Algebra and its Applications*, 433(2): 401–411, 2010.
9. K. Benson, D. Ferrero, M. Flagg, V. Furst, L. Hogben, V. Vasilevska, and B. Wissman. Power domination and zero forcing. *arXiv*:1510.02421, 2015.
10. A. Berliner, C. Bozeman, S. Butler, M. Catral, L. Hogben, B. Kroschel, J.C-H. Lin, N. Warnberg, M. Young. Zero forcing propagation time on oriented graphs. *Submitted*, 2015.
11. B. Brimkov. Complexity and computation of connected zero forcing. *arXiv*:1607.00658, 2016.
12. B. Brimkov and R. Davila. Characterizations of the connected forcing number of a graph. *arXiv*:1604.00740, 2016.
13. B. Brimkov, C.C. Fast, I.V. Hicks. Graphs with extremal connected forcing numbers. *arXiv*:1701.08500, 2017.
14. A. Buchanan, J.S. Sung, S. Butenko, and E.L. Pasiliao. An integer programming approach for fault-tolerant connected dominating sets. *INFORMS Journal on Computing*, 27(1): 178–188, 2015.
15. D. Burgarth and V. Giovannetti. Full control by locally induced relaxation. *Physical Review Letters*, 99(10): 100501, 2007.
16. S. Butler, and M. Young. Throttling zero forcing propagation speed on graphs. *Australasian Journal of Combinatorics*, 57: 65–71, 2013.
17. D. Burgarth, V. Giovannetti, L. Hogben, S. Severini, and M. Young. Logic circuits from zero forcing. *arXiv*:1106.4403, 2011.
18. S. Butler, L. DeLoss, J. Grout, H.T. Hall, J. LaGrange, T. McKay, J. Smith, and G. Tims. Minimum Rank Library (*Sage* programs for calculating bounds on the minimum rank of a graph, and for computing zero forcing parameters), 2014. Available at `https://github.com/jasongrout/minimum_rank`.
19. S. Butler, J. Grout, and H.T. Hall. Using variants of zero forcing to bound the inertia set of a graph. *Electronic Journal of Linear Algebra*, 30: 2015.
20. Y. Caro, D.B. West, and R. Yuster. Connected domination and spanning trees with many leaves. *SIAM Journal on Discrete Mathematics*, 13: 202–211, 2000.
21. R. Carvajal, M. Constantino, M. Goycoolea, J.P. Vielma, and A. Weintraub. Imposing connectivity constraints in forest planning models. *Operations Research*, 61(4): 824–836, 2013.
22. K.B. Chilakamarri, N. Dean, C.X. Kang, and E. Yi. Iteration index of a zero forcing set in a graph. *arXiv*:1105.1492, 2011.
23. G. Dantzig, R. Fulkerson, and S. Johnson. Solution of a large-scale traveling-salesman problem. *Operations Research*, 2 393–410, 1954.
24. W.J. Desormeaux, T.W. Haynes, and M.A. Henning. Bounds on the connected domination number of a graph. *Discrete Applied Mathematics*, 161(18): 2925–2931, 2013.
25. C. Edholm, L. Hogben, J. LaGrange, and D. Row. Vertex and edge spread of zero forcing number, maximum nullity, and minimum rank of a graph. *Linear Algebra and its Applications*, 436(12): 4352–4372, 2012.
26. J. Ekstrand, et al. Positive semidefinite zero forcing. *Linear Algebra and its Applications*, 439(7): 1862–1874, 2013.

27. L. Eroh, C. Kang, and E. Yi. Metric dimension and zero forcing number of two families of line graphs. *arXiv*:1207.6127, 2012.
28. S. Fallat and L. Hogben. The minimum rank of symmetric matrices described by a graph: A survey. *Linear Algebra and its Applications*, 426: 558–582, 2007.
29. N. Fan and J.-P. Watson. Solving the connected dominating set problem and power dominating set problem by integer programming. *International Conference on Combinatorial Optimization and Applications*, Springer Berlin Heidelberg, pp. 371–383, 2012.
30. C.C. Fast and I.V. Hicks. The effect of vertex degrees on the zero-forcing number and iteration index of a graph. *Submitted*, 2016.
31. M. Fischetti, M. Leitner, I. Ljubić, M. Luipersbeck, M. Monaci, M. Resch, D. Salvagnin, and M. Sinnl. Thinning out Steiner trees: a node-based model for uniform edge costs. *Mathematical Programming Computation*, 1–27, 2016.
32. M. Fischetti and A. Lodi. Optimizing over the first Chvátal closure. *Mathematical Programming*, 110: 3–20, 2007.
33. F.V. Fomin, F. Grandoni, and D. Kratsch. Solving connected dominating set faster than 2^n. *Algorithmica* 52(2): 153–166, 2008.
34. F. Goldberg and A. Berman. Zero forcing for sign patterns. *Linear Algebra and its Applications*, 447: 56–67, 2014.
35. T. Haynes, S. Hedetniemi, S. Hedetniemi, and M. Henning. Domination in graphs applied to electric power networks. *SIAM Journal on Discrete Mathematics*, 15(4): 519–529, 2002.
36. L. Hogben, N. Kingsley, S. Meyer, S. Walker, and M. Young. Propagation time for zero forcing on a graph. *Discrete Applied Mathematics*, 160(13): 1994–2005, 2012.
37. L. Hogben, K. F. Palmowski, D. E. Roberson, and M. Young. Fractional zero forcing via three-color forcing games. *Discrete Applied Mathematics*, 213: 114-129, 2016.
38. L.-H. Huang, G. J. Chang, and H.-G. Yeh. On minimum rank and zero forcing sets of a graph. *Linear Algebra and its Applications*, 432: 2961–2973, 2010
39. L. Lu, B. Wu, and Z. Tang. Proof of a conjecture on the zero forcing number of a graph. *Discrete Applied Mathematics*, 213: 223-237, 2016.
40. S. Meyer. Zero forcing sets and bipartite circulants. *Linear Algebra and its Applications*, 436(4): 888–900, 2012.
41. C.E. Miller, A.W. Tucker, and R.A. Zemlin. Integer programming formulation of traveling salesman problems. *Journal of the ACM*, 7: 326–329, 1960.
42. F.P. Quintão, A.S. da Cunha, G.R. Mateus, and A. Lucena. The k-cardinality tree problem: reformulations and Lagrangian relaxation. *Discrete Applied Mathematics*, 158(12): 1305 – 1314, 2010.
43. D.D. Row. A technique for computing the zero forcing number of a graph with a cut-vertex. *Linear Algebra and its Applications*, 436: 4423–4432, 2012.
44. M. Trefois and J.C. Delvenne. Zero forcing number, constrained matchings and strong structural controllability. *arXiv*:1405.6222v2, 2015.
45. Y. Wang, A. Buchanan, and S. Butenko. On imposing connectivity constraints in integer programs. *Optimization Online*, 2015.
46. N. Warnberg. Positive semidefinite propagation time. *Discrete Applied Mathematics*, 198: 274–290, 2016.
47. D.J. Watts and S.H. Strogatz. Collective dynamics of 'small-world' networks. *Nature*, 393: 440–442, 1998.
48. D.B. West. *Introduction to Graph Theory*. Prentice Hall, Inc., Upper Saddle River, NJ, 2001.

49. B. Yang. Fast-mixed searching and related problems on graphs. *Theoretical Computer Science*, 507: 100–113, 2013.
50. M. Zhao, L. Kang, and G. Chang. Power domination in graphs. *Discrete Mathematics*, 306(15): 1812–1816, 2006.